

Folding Domain-Specific Languages: Deep and Shallow Embeddings

(Functional Pearl)

Jeremy Gibbons Nicolas Wu

Department of Computer Science, University of Oxford
{jeremy.gibbons,nicolas.wu}@cs.ox.ac.uk

Abstract

A domain-specific language can be implemented by embedding within a general-purpose host language. This embedding may be *deep* or *shallow*, depending on whether terms in the language construct syntactic or semantic representations. The deep and shallow styles are closely related, and intimately connected to folds; in this paper, we explore that connection.

1. Introduction

General-purpose programming languages (GPLs) are great for generality. But this very generality can count against them: it may take a lot of programming to establish a suitable context for a particular domain; and the programmer may end up being spoilt for choice with the options available to her—especially if she is a domain specialist rather than primarily a software engineer. This tension motivates many years of work on techniques to support the development of *domain-specific languages* (DSLs) such as VHDL, SQL and PostScript: languages specialized for a particular domain, incorporating the contextual assumptions of that domain and guiding the programmer specifically towards programs suitable for that domain.

There are two main approaches to DSLs. *Standalone* DSLs provide their own custom syntax and semantics, and standard compilation techniques are used to translate or interpret programs written in the DSL for execution. Standalone DSLs can be designed for maximal convenience to their intended users. But the exercise can be a significant undertaking for the implementer, involving an entirely separate ecosystem—compiler, editor, debugger, and so on—and typically also much reinvention of standard language features such as local definitions, conditionals, and iteration.

The alternative approach is to *embed* the DSL within a host GPL, essentially as a collection of definitions written in the host language. All the existing facilities and infrastructure of the host environment can be appropriated for the DSL, and familiarity with the syntactic conventions and tools of the host language can be carried over to the DSL. Whereas the standalone approach is the

most common one within object-oriented circles [10], the embedded approach is typically favoured by functional programmers [18]. It seems that core FP features such as algebraic datatypes and higher-order functions are extremely helpful in defining embedded DSLs; conversely, it has been said [23] that language-oriented tasks such as DSLs are the killer application for FP.

Amongst embedded DSLs, there are two further refinements. With a *deep embedding*, terms in the DSL are implemented simply to construct an abstract syntax tree (AST), which is subsequently transformed for optimization and traversed for evaluation. With a *shallow embedding*, terms in the DSL are implemented directly by their semantics, bypassing the intermediate AST and its traversal. The names ‘deep’ and ‘shallow’ seem to have originated in the work of Boulton and colleagues on embedding hardware description languages in theorem provers for the purposes of verification [6]. Boulton’s motivation for the names was that a deep embedding preserves the syntactic representation of a term, “whereas in a shallow embedding [the syntax] is just a surface layer that is easily blown away by rewriting” [5]. It turns out that deep and shallow embeddings are closely related, and intimately connected to folds; our purpose in this paper is to explore that connection.

2. Embedding DSLs

We start by looking a little closer at deep and shallow embeddings. Consider a very simple language of arithmetic expressions, involving integer constants and addition:

```
type Expr1 = ...  
lit  :: Integer      → Expr1  
add  :: Expr1 → Expr1 → Expr1
```

The expression $(3 + 4) + 5$ is represented in the DSL by the term `add (add (lit 3) (lit 4)) (lit 5)`.

As a deeply embedded DSL, the two operations `lit` and `add` are encoded directly as constructors of an algebraic datatype:

```
data Expr2 :: * where  
  Lit  :: Integer      → Expr2  
  Add  :: Expr2 → Expr2 → Expr2  
  
lit n    = Lit n  
add x y  = Add x y
```

(We have used Haskell’s ‘generalized algebraic datatype’ notation, in order to make the types of the constructors `Lit` and `Add` explicit; but we are not using the generality of GADTs here, and the old-fashioned way would have worked too.) Observations of terms in the DSL are defined as functions over the algebraic datatype. For example, here is how to evaluate an expression:

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ICFP '14, September 1–6, 2014, Gothenburg, Sweden.
Copyright © 2014 ACM 978-1-4503-2873-9/14/09...\$15.00.
<http://dx.doi.org/10.1145/2628136.2628138>




Figure 1. The Brent-Kung parallel prefix circuit of width 4

```
eval2 :: Expr2 → Integer
eval2 (Lit n)    = n
eval2 (Add x y) = eval2 x + eval2 y
```

This might be used as follows:

```
> eval2 (Add (Add (Lit 3) (Lit 4)) (Lit 5))
12
```

In other words, a deep embedding consists of a representation of the abstract syntax as an algebraic datatype, together with some functions that assign semantics to that syntax by traversing the algebraic datatype.

A shallow embedding eschews the algebraic datatype, and hence the explicit representation of the abstract syntax of the language; instead, the language is defined directly in terms of its semantics. For example, if the semantics is to be evaluation, then we could define:

```
type Expr3 = Integer
lit n      = n
add x y    = x + y
eval3 :: Expr3 → Integer
eval3 n    = n
```

This might be used as follows:

```
> eval3 (add (add (lit 3) (lit 4)) (lit 5))
12
```

We have used subscripts to distinguish different representations of morally ‘the same’ functions ($eval_2$ and $eval_3$) and types ($Expr_2$ and $Expr_3$). We will continue that convention throughout the paper.

One might see the deep and shallow embeddings as duals, in a variety of senses. For one sense, the language constructs *Lit* and *Add* in the deep embedding do none of the work, leaving this entirely to the observation function *eval*; in contrast, in the shallow embedding, the language constructs *lit* and *add* do all the work, and the observer *eval₃* is simply the identity function.

For a second sense, it is trivial to add a second observer such as pretty-printing to the deep embedding—just define another function alongside *eval*—but awkward to add a new construct such as multiplication: doing so entails revisiting the definitions of all existing observers to add an additional clause. In contrast, adding a construct to the shallow embedding—alongside *lit* and *add*—is trivial, but the obvious way of introducing an additional observer entails completely revising the semantics by changing the definitions of all existing constructs. This is precisely the tension underlying the *expression problem* [22, 29], so named for precisely this example.

The types of *lit* and *add* in the shallow embedding coincide with those of *Lit* and *Add* in the deep embedding; moreover, the definitions of *lit* and *add* in the shallow embedding correspond to the ‘actions’ in each clause of the definition of the observer in the deep embedding. The shallow embedding presents a *compositional* semantics for the language, since the semantics of a composite term is explicitly composed from the semantics of its components. Indeed, it is only such compositional semantics that can be captured in a shallow embedding; it is possible to define a more sophisticated non-



Figure 2. Identity circuit *identity 4* and fan circuit *fan 4* of width 4

compositional semantics as an interpretation of a deep embedding, but not possible to represent that semantics directly via a shallow embedding.

However, there is no duality in the categorical sense of reversing arrows. Although deep and shallow embeddings have been called the ‘initial’ and ‘final’ approaches [8], in fact the two approaches are equivalent, and both correspond to initial algebras; Carette *et al.* say only that they use the term ‘final’ “because we represent each object term not by its abstract syntax but by its denotation in a semantic algebra”, and they are not concerned with final coalgebras.

3. Scans

The expression language above is very simple—perhaps too simple to serve as a convincing vehicle for discussion. As a more interesting example of a DSL, we turn to a language for parallel prefix circuits [13], which crop up in a number of different applications—carry-lookahead adders, parallel sorting, and stream compaction, to name but a few. Given an associative binary operator \bullet , a prefix computation of width $n > 0$ takes a sequence x_1, x_2, \dots, x_n of inputs and produces the sequence $x_1, x_1 \bullet x_2, \dots, x_1 \bullet x_2 \bullet \dots \bullet x_n$ of outputs. A parallel prefix circuit performs this computation in parallel, in a fixed format independent of the input values x_i .

An example of such a circuit is depicted in Figure 1. This circuit diagram should be read as follows. The inputs are fed in at the top, and the outputs fall out at the bottom. Each node (the blobs in the diagram) represents a local computation, combining the values on each of its input wires using \bullet , in left-to-right order, and providing copies of the result on each of its output wires. It is an instructive exercise to check that this circuit does indeed take x_1, x_2, x_3, x_4 to $x_1, x_1 \bullet x_2, x_1 \bullet x_2 \bullet x_3, x_1 \bullet x_2 \bullet x_3 \bullet x_4$.

Such circuits can be constructed using the following operators:

```
type Size = Int -- positive
type Circuit1 = ...
identity :: Size → Circuit1
fan      :: Size → Circuit1
above   :: Circuit1 → Circuit1 → Circuit1
beside   :: Circuit1 → Circuit1 → Circuit1
stretch  :: [Size] → Circuit1 → Circuit1
```

The most basic building block is the identity circuit, *identity n*, which creates a circuit consisting of n parallel wires that copy input to output. The other primitive is the fan circuit; *fan n* takes n inputs, and adds its first input to each of the others. We only consider non-empty circuits, so n must be positive in both cases. Instances of *identity* and *fan* of width 4 are shown in Figure 2.

Then there are three combinators for circuits. The series or vertical composition, *above c d*, takes two circuits c and d of the same width, and connects the outputs of c to the inputs of d . The parallel or horizontal composition, *beside c d*, places c beside d , leaving them unconnected; there are no width constraints on c and d . Figure 3 shows a 2-fan beside a 1-identity, a 1-identity beside a 2-fan, and the first of these above the second (note that they both have width 3); this yields the “serial” parallel prefix circuit of width 3.

Finally, the stretch combinator, *stretch ws c*, takes a non-empty list of positive widths $ws = [w_1, \dots, w_n]$ of length n , and a circuit c of width n , and “stretches” c out to width $\text{sum } ws$ by interleaving some additional wires. Of the first bundle of w_1 inputs, the last is routed to the first input of c and the rest pass straight through; of




Figure 3. The construction of a parallel prefix circuit of width 3

the next bundle of w_2 inputs, the last is routed to the second input of c and the rest pass straight through; and so on. (Note that each bundle width w_i must be positive.) For example, Figure 4 shows a 3-fan stretched out to width 8, in bundles of $[3, 2, 3]$.

So one possible construction of the Brent–Kung parallel prefix circuit in Figure 1 is

```
(fan 2 'beside' fan 2) 'above'
stretch [2, 2] (fan 2) 'above'
(identity 1 'beside' fan 2 'beside' identity 1)
```

The general Brent–Kung construction [7] is given recursively. The general pattern is a row of 2-fans, possibly with an extra wire in the case of odd width; then a Brent–Kung circuit of half the width, stretched out by a factor of two; then another row of 2-fans, shifted one place to the right.

```
brentkung :: Size → Circuit1
brentkung 1 = identity 1
brentkung w
  = (row (replicate u (fan 2)) 'pad' w) 'above'
    (stretch (replicate u 2) (brentkung u) 'pad' w) 'above'
    (row (identity 1 : replicate v (fan 2)) 'pad' (w - 1))
  where (u, v) = (w `div` 2, (w - 1) `div` 2)
        c 'pad' w = if even w then c else c 'beside' identity 1
        row      = foldr1 beside
```

The Brent–Kung circuit of width 16 is shown in Figure 5. Note one major benefit of defining *Circuit* as an embedded rather than a standalone DSL: we can exploit for free host language constructions such as *replicate* and *foldr1*, rather than having to reinvent them within the DSL.

As a deeply embedded DSL, circuits can be captured by the following algebraic datatype:

```
data Circuit2 :: * where
  Identity :: Size → Circuit2
  Fan      :: Size → Circuit2
  Above    :: Circuit2 → Circuit2 → Circuit2
  Beside   :: Circuit2 → Circuit2 → Circuit2
  Stretch  :: [Size] → Circuit2 → Circuit2
```

It is, of course, straightforward to define functions to manipulate this representation. Here is one, which computes the width of a circuit:

```
type Width = Int
width2 :: Circuit2 → Width
width2 (Identity w) = w
width2 (Fan w)      = w
width2 (Above x y)  = width2 x
width2 (Beside x y) = width2 x + width2 y
width2 (Stretch ws x) = sum ws
```

Note that *width₂* is compositional: it is a fold over the abstract syntax of *Circuit₂*s. That makes it a suitable semantics for a shallow embedding. That is, we could represent circuits directly by their widths, as follows:

```
type Circuit3 = Width
identity w = w
```




Figure 4. A 3-fan stretched out by widths $[3, 2, 3]$

```
fan w      = w
above x y  = x
beside x y = x + y
stretch ws x = sum ws
width3 :: Circuit3 → Width
width3 = id
```

Clearly, width is a rather uninteresting semantics to give to circuits. But what other kinds of semantics will fit the pattern of compositionality, and so be suitable for a shallow embedding? In order to explore that question, we need to look a bit more closely at folds and their variations.

4. Folds

Folds are the natural pattern of computation induced by inductively defined algebraic datatypes. We consider here just polynomial algebraic datatypes, namely those with one or more constructors, each constructor taking zero or more arguments to the datatype being defined, and each argument either having a fixed type independent of the datatype, or being a recursive occurrence of the datatype itself. For example, the polynomial algebraic datatype *Circuit₂* above has five constructors; *Identity* and *Fan* each take one argument of the fixed type *Size*; *Above* and *Beside* take two arguments, both recursive occurrences; *Stretch* takes two arguments, one of which is the fixed type *[Size]*, and the other is a recursive argument. Thus, we rule out contravariant recursion, polymorphic datatypes, higher kinds, and other such esoterica. For simplicity, we also ignore DSLs with binding constructs, which complicate matters significantly; for more on this, see [1, 8].

The general case is captured by a shape—also called a base or pattern functor—which is an instance of the *Functor* type class:

```
class Functor f where
  fmap :: (a → b) → (f a → f b)
```

For *Circuit₂*, the shape is given by *CircuitF* as follows, where the parameter *x* marks the recursive spots:

```
data CircuitF :: * → * where
  IdentityF :: Size → CircuitF x
  FanF      :: Size → CircuitF x
  AboveF    :: x → x → CircuitF x
  BesideF   :: x → x → CircuitF x
  StretchF  :: [Size] → x → CircuitF x

instance Functor CircuitF where
  fmap f (IdentityF w) = IdentityF w
  fmap f (FanF w)      = FanF w
  fmap f (AboveF x1 x2) = AboveF (f x1) (f x2)
  fmap f (BesideF x1 x2) = BesideF (f x1) (f x2)
  fmap f (StretchF ws x) = StretchF ws (f x)
```

We can use this shape functor as the basis of an alternative definition of the algebraic datatype *Circuit₂*:

```
data Circuit4 = In (CircuitF Circuit4)
```

Now, an algebra for a functor *f* consists of a type *a* and a function taking an *f*-structure of *a*-values to an *a*-value. For the functor *CircuitF*, this is:

```
type CircuitAlg a = CircuitF a → a
```




Figure 5. The Brent–Kung parallel prefix circuit of width 16

Such an algebra is precisely the information needed to fold a data structure:

```
foldC :: CircuitAlg a → Circuit4 → a
foldC h (In x) = h (fmap (foldC h) x)
```

For example, *width* is a fold for the deeply embedded DSL of shape *CircuitF*, and is determined by the following algebra:

```
widthAlg :: CircuitAlg Width
widthAlg (IdentityF w) = w
widthAlg (FanF w)      = w
widthAlg (AboveF x y)  = x
widthAlg (BesideF x y) = x + y
widthAlg (StretchF ws x) = sum ws

width4 :: Circuit4 → Width
width4 = foldC widthAlg
```

So a compositional observation function for the deep embedding, such as *width₄*, is precisely a fold using such an algebra. We know a lot about folds, and this tells us a lot about embedded DSLs. We discuss these consequences next.

4.1 Multiple interpretations

As mentioned above, the deep embedding smoothly supports additional observations. For example, suppose that we also wanted to find the depth of our circuits. No problem—we can just define another observation function.

```
type Depth = Int

depthAlg :: CircuitAlg Depth
depthAlg (IdentityF w) = 0
depthAlg (FanF w)      = 1
depthAlg (AboveF x y)  = x + y
depthAlg (BesideF x y) = x `max` y
depthAlg (StretchF ws x) = x

depth4 :: Circuit4 → Depth
depth4 = foldC depthAlg
```

But what about with a shallow embedding? With this approach, circuits can only have a single semantics, so how do we accommodate finding both the width and the depth of a circuit? It's not much more difficult than with a deep embedding; we simply make the semantics a pair, providing both interpretations simultaneously.

```
type Circuit5 = (Width, Depth)
```

Now the observation functions *width₅* and *depth₅* become projections, rather than just the identity function.

```
width5 :: Circuit5 → Width
width5 = fst

depth5 :: Circuit5 → Depth
depth5 = snd
```

The individual operations can be defined much as before, just by projecting the relevant components out of the pair:

```
wdAlg :: CircuitAlg Circuit5
wdAlg (IdentityF w) = (w, 0)
wdAlg (FanF w)      = (w, 1)
wdAlg (AboveF x y)  = (width5 x, depth5 x + depth5 y)
wdAlg (BesideF x y) = (width5 x + width5 y,
                      depth5 x `max` depth5 y)
wdAlg (StretchF ws x) = (sum ws, depth5 x)
```

This algebra is the essence of the shallow embedding; for example,

```
identity5 w = wdAlg (IdentityF w)
```

and so on. Of course, this works better under lazy than under eager evaluation: if only one of the two interpretations of an expression is needed, only that one is evaluated. And it's rather clumsy from a modularity perspective; we will return to this point later.

Seen from the fold perspective, this step is no surprise: the ‘banana split law’ [9] tells us that tupling two independent folds gives another fold, so multiple interpretations can be provided in the shallow embedding nearly as easily as in the deep embedding.

4.2 Dependent interpretations

A shallow embedding supports only compositional interpretations, whereas a deep embedding provides full access to the AST and hence also non-compositional manipulations. Here, ‘compositionality’ of an interpretation means that the interpretation of a whole may be determined solely from the interpretations of its parts; it is both a valuable property for reasoning and a significant limitation to expressivity. Not all interpretations are of this form; sometimes a ‘primary’ interpretation of the whole depends also on ‘secondary’ interpretations of its parts.

For example, whether a circuit is well formed depends on the widths of its constituent parts. Given that we have an untyped (or rather, ‘unsized’) model of circuits, we might capture this property in a separate function *wellSized*:

```
type WellSized = Bool

wellSized :: Circuit2 → WellSized
wellSized (Identity w) = True
wellSized (Fan w)      = True
```

$$\begin{aligned}
\text{wellSized } (\text{Above } x \ y) &= \text{wellSized } x \wedge \text{wellSized } y \\
&\quad \wedge \text{width } x \equiv \text{width } y \\
\text{wellSized } (\text{Beside } x \ y) &= \text{wellSized } x \wedge \text{wellSized } y \\
\text{wellSized } (\text{Stretch } ws \ x) &= \text{wellSized } x \wedge \text{length } ws \equiv \text{width } x
\end{aligned}$$

This is a non-compositional interpretation of the abstract syntax, because *wellSized* sometimes depends on the *width* of subcircuits as well as their recursive image under *wellSized*. In other words, *wellSized* is not a fold, and there is no corresponding *CircuitAlg*.

What can we do about such non-compositional interpretations in the shallow embedding? Again, fold theory comes to the rescue: *wellSized* and *width* together form a *mutumorphism* [9]—that is, two mutually dependent folds—and the tuple of these two functions again forms a fold. (In fact, this is a special case, a *zygomorphism* [9], since the dependency is only one-way. Simpler still, we have seen another special case in the banana split above, where neither of the two folds depends on the other.)

```

type Circuit6 = (WellSized, Width)
wswAlg :: CircuitAlg Circuit6
wswAlg (IdentityF w) = (True, w)
wswAlg (FanF w)      = (True, w)
wswAlg (AboveF x y)  = (fst x ∧ fst y ∧ snd x ≡ snd y, snd x)
wswAlg (BesideF x y) = (fst x ∧ fst y, snd x + snd y)
wswAlg (StretchF ws x) = (fst x ∧ length ws ≡ snd x, sum ws)

```

So although *wellSized* = *fst* ∘ *foldC* *wswAlg* is not a fold, it is manifestly clear that *foldC* *wswAlg* is. Tupling functions in this way is analogous to strengthening the invariant of an imperative loop to record additional information [19], and is a standard trick in program calculation [17].

Another example of a dependent interpretation is provided by what Hinze [13] calls the *standard model* of the circuit: its interpretation as a computation. As discussed in the introduction, this is defined in terms of an associative binary operator (•), which we capture by the following type class:

```

class Semigroup s where
  (•) :: s → s → s -- • is associative

```

The interpretation *apply* interprets a circuit of width *n* as a function operating on lists of length *n*:

```

apply :: Semigroup a ⇒ Circuit2 → [a] → [a]
apply (IdentityF w) xs = xs
apply (FanF w) (x : xs) = x : map (x •) xs
apply (Above c d) xs = apply d (apply c xs)
apply (Beside c d) xs = apply c xs ++ apply d xs
where (ys, zs) = splitAt (width c) xs
apply (Stretch ws c) xs = concat
  (zipWith snoc (map init xs) (apply c (map last xs)))
where xs = bundle ws xs

```

Here, *snoc* is ‘cons’ backwards,

```
snoc ys z = ys ++ [z]
```

and *bundle* *ws* *xs* groups the list *xs* into bundles of widths *ws*, assuming that *sum* *ws* ≡ *length* *xs*:

```

bundle :: Integral i ⇒ [i] → [a] → [[a]]
bundle [] [] = []
bundle (w : ws) xs = ys : bundle ws zs
where (ys, zs) = splitAt w xs

```

The *apply* interpretation is another zygomorphism, because in the *Beside* case *apply* depends on *width* *c* as well as *apply* *c* and *apply* *d*. And indeed, Hinze’s ‘standard model’ [13] comprises both the list transformer and the width, tupled together.

4.3 Context-sensitive interpretations

Consider generating a circuit layout from a circuit description, for example as the first step in expressing the circuit in a hardware description language such as VHDL—or, for that matter, for producing the diagrams in this paper. The essence of the translation is to determine the connections between vertical wires. Note that each circuit can be thought of as a sequence of layers, and connections only go from one layer to the next (and only rightwards, too). So it suffices to generate a list of layers, where each layer is a collection of pairs (i, j) denoting a connection from wire *i* on this layer to wire *j* on the next. The ordering of the pairs on each layer is not significant. We count from 0. For example, the Brent–Kung circuit of size 4 given in Figure 1 has the following connections:

$$[[(0, 1), (2, 3)], [(1, 3)], [(1, 2)]]$$

That is, there are three layers; the first layer has connections from wire 0 to wire 1 and from wire 2 to wire 3; the second a single connection from wire 1 to wire 3; and the third a single connection from wire 1 to wire 2.

```

type Layout = [[(Size, Size)]]
layout :: Circuit2 → Layout
layout (IdentityF w) = []
layout (FanF w)      = [ [ (0, j) | j ← [1..w-1] ] ]
layout (Above c d)   = layout c ++ layout d
layout (Beside c d)   = lzw (++) (layout c)
                        (shift (width c) (layout d))
layout (Stretch ws c) = map (map (connect ws)) (layout c)
shift w               = map (map (pmap (w+)))
connect ws             = pmap (pred ∘ ((scanl1 (+) ws)!!))

```

Here, *pmap* is the map function for homogeneous pairs:

```

pmap :: (a → b) → (a, a) → (b, b)
pmap f (x, y) = (f x, f y)

```

The function *lzw* is ‘long zip with’ [12], which zips two lists together and returns a result as long as the longer argument. The binary operator is used to combine corresponding elements; if one list is shorter then the remaining elements of the other are simply copied.

```

lzw :: (a → a → a) → [a] → [a] → [a]
lzw f [] ys = ys
lzw f xs [] = xs
lzw f (x : xs) (y : ys) = f x y : lzw f xs ys

```

The *layout* interpretation is yet another zygomorphism, because *layout* (*Beside* *c* *d*) depends on *width* *c* as well as *layout* *c* and *layout* *d*. In fact, in general we need the width of the circuit anyway in order to determine the layout, in case the rightmost wire is not connected to the others. So the techniques discussed above will allow us to express the layout as a shallow embedding, whose essence is as follows:

```

lwAlg :: CircuitAlg (Layout, Width)
lwAlg (IdentityF w) = ([], w)
lwAlg (FanF w)      = ([ [ (0, j) | j ← [1..w-1] ] ], w)
lwAlg (AboveF c d)  = (l1 ++ l2, w2)
where (l1, w1) = c; (l2, w2) = d
lwAlg (BesideF c d) = (lzw (++) l1 (shift w1 l2), w1 + w2)
where (l1, w1) = c; (l2, w2) = d
lwAlg (StretchF ws (l, w)) = (map (map (connect ws)) l, sum ws)

```

But even having achieved this, there is room for improvement. In the *Beside* and *Stretch* clauses, sublayouts are postprocessed using *shift* and *map* (*map* (*connect* *ws*)) respectively. It would be more efficient to do this processing via an *accumulating parameter* [3]

instead. In this case, a transformation on wire indices suffices as the accumulating parameter (*‘tlayout’* stands for ‘transformed layout’):

```
tlayout :: (Size → Size) → Circuit2 → Layout
tlayout f c = map (map (pmap f)) (layout c)
```

Of course, *layout* = *tlayout id*, and it is a straightforward exercise to synthesize the following more efficient definition of *tlayout*:

```
tlayout :: (Size → Size) → Circuit2 → Layout
tlayout f (Identity w) = []
tlayout f (FanF w) = [[(f 0, f j) | j ← [1..w-1]]]
tlayout f (Above c d) = tlayout f c ++ tlayout f d
tlayout f (Beside c d) = lzw (++) (tlayout f c)
                                   (tlayout ((w+) ∘ f) d)
```

```
where w = width c
tlayout f (Stretch ws c) = tlayout (pred ∘ (vs!!) ∘ f) c
where vs = scanl1 (+) ws
```

And how does this work out with a shallow embedding? Note that *tlayout f* is no longer a fold, because the accumulating parameter changes in some recursive calls. One might say that *tlayout* is a context-sensitive layout function, and the context may vary in recursive calls. But standard fold technology comes to the rescue once more: *tlayout* may not be a fold, but *flip tlayout* is—specifically, an accumulating fold.

```
tlwAlg :: CircuitAlg ((Size → Size) → Layout, Width)
tlwAlg (IdentityF w) = (λf → [], w)
tlwAlg (FanF w) = (λf → [(f 0, f j) | j ← [1..w-1]], w)
tlwAlg (AboveF c d) = (λf → fst c f ++ fst d f, snd c)
tlwAlg (BesideF c d) = (λf → lzw (++) (fst c f)
                                   (fst d ((snd c +) ∘ f)),
                                   snd c + snd d)
tlwAlg (StretchF ws c) = (λf → fst c (pred ∘ (vs!!) ∘ f), sum ws)
where vs = scanl1 (+) ws
```

The alert reader may have noted another source of inefficiency in *layout*, namely the uses of *++* and *lzw (++)* in the *Above* and *Beside* cases. These too can be removed, by introducing two more accumulating parameters, giving:

```
ulayout :: (Size → Size) → Layout → Layout →
Circuit2 → Layout
ulayout f l' c = (lzw (++) (map (map (pmap f)) (layout c)) l) ++ l'
```

(now *‘ulayout’* stands for ‘ultimate layout’). From this specification we can synthesize a definition that takes linear time in the ‘size’ of the circuit, for a reasonable definition of ‘size’. We leave the details as an exercise.

In fact, the standard interpretation *apply* given above is really another accumulating fold, in disguise. Rather than reading the type

```
apply :: Semigroup a ⇒ Circuit2 → [a] → [a]
```

as defining an interpretation of circuits as list transformers of type *Semigroup a ⇒ ([a] → [a])*, one can read it as defining a context-dependent interpretation as an output list of type *Semigroup a ⇒ [a]*, dependent on some input list of the same type. The interpretation is implemented in terms of an accumulating parameter; this is initially the input list, but it ‘accumulates’ by attrition via *splitAt* and *map last ∘ bundle ws* as the evaluation proceeds.

4.4 Parametrized interpretations

We saw in Section 4.1 that it is not difficult to provide multiple interpretations with a shallow embedding, by constructing a tuple as the semantics of an expression and projecting the desired interpretation from the tuple. But this is still a bit clumsy: it entails revising existing code each time a new interpretation is added, and wide tuples generally lack good language support [24].

But as we have also seen, all compositional interpretations conform to a common pattern: they are folds. So we can provide a shallow embedding as precisely that pattern—that is, in terms of a single *parametrized* interpretation, which is a higher-order value representing the fold.

```
newtype Circuitγ = Cγ { unCγ :: ∀a. CircuitAlg a → a }
identityγ w = Cγ (λh → h (IdentityF w))
fanγ w = Cγ (λh → h (FanF w))
aboveγ x y = Cγ (λh → h (AboveF (unCγ x h) (unCγ y h)))
besideγ x y = Cγ (λh → h (BesideF (unCγ x h) (unCγ y h)))
stretchγ ws x = Cγ (λh → h (StretchF ws (unCγ x h)))
```

(We need the **newtype** instead of a plain **type** synonym because of the quantified type.) This shallow encoding subsumes all others; it specializes to *depth* and *width*, and of course to any other fold:

```
widthγ :: Circuitγ → Width
widthγ circuit = unCγ circuit widthAlg
depthγ :: Circuitγ → Depth
depthγ circuit = unCγ circuit depthAlg
```

In fact, the shallow embedding provides a universal generic interpretation as the *Church encoding* [14] of the AST—or more precisely, because it is typed, the *Böhm–Berarducci encoding* [4].

Universality is witnessed by the observation that it is possible to recover the deep embedding from this one ‘mother of all shallow embeddings’ [8]:

```
deep :: Circuitγ → Circuit4
deep circuit = unCγ circuit In
```

(So it turns out that the syntax of the DSL is not really as ephemeral in a shallow embedding as Boulton’s choice of terms [6] suggests.) And conversely, one can map from the deep embedding to the parametrized shallow embedding, and thence to any other shallow embedding:

```
shallow :: Circuit4 → Circuitγ
shallow = foldC shallowAlg
shallowAlg :: CircuitAlg Circuitγ
shallowAlg (IdentityF w) = identityγ w
shallowAlg (FanF w) = fanγ w
shallowAlg (AboveF c d) = aboveγ c d
shallowAlg (BesideF c d) = besideγ c d
shallowAlg (StretchF ws c) = stretchγ ws c
```

Moreover, *deep* and *shallow* are each other’s inverses, assuming parametricity [28].

4.5 Implicitly parametrized interpretations

The shallow embedding in Section 4.4 involves explicitly passing an algebra with which to interpret terms. That parameter may be passed implicitly instead, if it can be determined from the type of the interpretation. In Haskell, this can be done by defining a suitable type class:

```
class Circuit8 circuit where
  identity8 :: Size → circuit
  fan8 :: Size → circuit
  above8 :: circuit → circuit → circuit
  beside8 :: circuit → circuit → circuit
  stretch8 :: [Size] → circuit → circuit
```

To specify a particular interpretation, one defines an instance of the type class for the type of that interpretation. For example, here is the specification of the ‘width’ interpretation:

```
newtype Width8 = Width { unWidth :: Int }
```

```

instance Circuit8 Width8 where
  identity8 w = Width w
  fan8 w      = Width w
  above8 x y  = x
  beside8 x y = Width (unWidth x + unWidth y)
  stretch8 ws x = Width (sum ws)

```

The **newtype** wrapper is often needed to allow multiple interpretations over the same underlying type; for example, we can provide both ‘width’ and ‘depth’ interpretations over integers:

```

newtype Depth8 = Depth { unDepth :: Int }
instance Circuit8 Depth8 where
  identity8 w = Depth 0
  fan8 w      = Depth 1
  above8 x y  = Depth (unDepth x + unDepth y)
  beside8 x y = Depth (unDepth x + max 1 unDepth y)
  stretch8 ws x = x

```

Some of the wrapping and unwrapping of *Width₈* and *Depth₈* values could be avoided by installing these types as instances of the *Num* and *Ord* type classes; this can even be done automatically in GHC, by exploiting the ‘Generalized Newtype Deriving’ extension.

The conventional implementation of type classes [30] involves constructing a *dictionary* for each type in the type class, and generating code that selects and passes the appropriate dictionary as an additional parameter to each overloaded member function (*identity₈*, *fan₈* etc). For an instance *c* of the type class *Circuit₈*, the dictionary type is equivalent to *CircuitAlg c*. Indeed, we might have defined instead

```

class Circuit9 c where
  alg :: CircuitAlg c
instance Circuit9 Width8 where
  alg = Width ∘ widthAlg ∘ fmap unWidth

```

so that the dictionary type is literally a *CircuitAlg c*: the Böhm-Berarducci and type-class approaches are really very similar.

4.6 Intermediate interpretations

Good practice in the design of embedded DSLs is to distinguish between a minimal ‘core’ language and a more useful ‘everyday’ language [26]. The former is more convenient for the language designer, but the latter more convenient for the language user. This apparent tension can be resolved by defining the additional constructs in the everyday language by translation to the core language.

For example, the *identity* construct in our DSL of circuits is redundant: *identity 1* is morally equivalent to *fan 1*, and for any other width *n*, we can construct a circuit equivalent to *identity n* by placing *n* copies of *identity 1* side by side (or alternatively, as *stretch [n] (identity 1)*). One might therefore identify a simpler datatype

```

data CoreCircuit :: * where
  CFan   :: Size → CoreCircuit
  CAbove :: CoreCircuit → CoreCircuit → CoreCircuit
  CBeside :: CoreCircuit → CoreCircuit → CoreCircuit
  CStretch :: [Size] → CoreCircuit → CoreCircuit

```

and use it as the carrier of a shallow embedding for the everyday language. The everyday constructs that correspond to core constructs are represented directly; the derived constructs are defined by translation.

```

type Circuit10 = CoreCircuit
coreAlg :: CircuitAlg Circuit10

```

```

coreAlg (IdentityF w) = foldr1 CBeside (replicate w (CFan 1))
coreAlg (FanF w)      = CFan w
coreAlg (AboveF x y)  = CAbove x y
coreAlg (BesideF x y) = CBeside x y
coreAlg (StretchF ws x) = CStretch ws x

```

One might see this as a shallow embedding, with the carrier *CoreCircuit* itself the deep embedding of a different, smaller language; the core constructs are implemented directly as constructors of *CoreCircuit*, and non-core constructs as a kind of ‘smart constructor’.

This suggests that ‘deep’ and ‘shallow’ do not form a dichotomy, but rather are two extreme points on a scale of embedding depth. Augustsson [2] discusses representations of intermediate depth, in which some constructs have deep embeddings and some shallow. In particular, for a language with a ‘semantics’ in the form of generated assembly code, the deeply embedded constructs will persist as generated code, whereas those with shallow embeddings will get translated away at ‘compile time’. Augustsson calls these *neritic* embeddings, after the region of the sea between the shore and the edge of the continental shelf.

4.7 Modular interpretations

The previous section explored cutting down the grammar of circuits by eliminating a constructor. Conversely, one might extend the grammar by adding constructors. Indeed, in addition to the ‘left stretch’ combinator we have used, Hinze [13] also provides a ‘right stretch’ combinator, which connects the first rather than the last wire of each bundle to the inner circuit. This is not needed in the core language, because it can be built out of existing components:

```

rstretch (ws ++ [w + 1]) c = stretch (1 : ws) c `beside` identity w

```

So one might extend the grammar of the everyday language, as embodied in the functor *CircuitF* or the type class *Circuit₈*, to incorporate this additional operator, but still use *CoreCircuit* as the actual representation.

Alternatively, one might hope for a modular technique for assembling embedded languages and their interpretations from parts, so that it is straightforward to add additional constructors like ‘right stretch’. Swierstra’s *datatypes à la carte* machinery [27] provides precisely such a thing, going some way towards addressing the expression problem discussed in Section 2.

The key idea is to represent each constructor separately:

```

data Identity11 c = Identity11 Size deriving Functor
data Fan11 c      = Fan11 Size deriving Functor
data Above11 c    = Above11 c c deriving Functor
data Beside11 c   = Beside11 c c deriving Functor
data Stretch11 c  = Stretch11 [Size] c deriving Functor

```

with a right-associating ‘sum’ operator for combining them:

```

data (f :+: g) e = Inl (f e) | Inr (g e) deriving Functor
infix :+:

```

One can assemble a functor from these components and make a deep embedding from it. For example, the sum of functors *CircuitF₁₁* is equivalent to *CircuitF* from the start of Section 4, and its fixpoint *Circuit₁₁* to *Circuit₄*:

```

type CircuitF11 = Identity11 :+: Fan11 :+: Above11 :+:
  Beside11 :+: Stretch11

```

```

data Fix f = In (f (Fix f))
type Circuit11 = Fix CircuitF11

```

This works, but it is rather clumsy. In particular, an expression of type *Circuit₁₁* involves a mess of *Inl*, *Inr* and *In* constructors, as seen in this rendition of the circuit in Figure 4:

```
stretchfan :: Circuit11
stretchfan = In (Inr (Inr (Inr (Inr (Stretch11 [3,2,3] (
  In (Inr (Inl (Fan11 3))))))))))
```

Fortunately, there is an obvious way of injecting payloads into sum types in this fashion, which we can express through a simple notion of subtyping between functors, witnessed by an injection:

```
class (Functor f, Functor g) => f <: g where
  inj :: f a -> g a
```

Subtyping is reflexive, and summands are subtypes of their sum:

```
instance Functor f => f <: f where
  inj = id

instance (Functor f, Functor g) => f <: (f :+: g) where
  inj = Inl

instance (Functor f, Functor g, Functor h, f <: g) =>
  f <: (h :+: g) where
  inj = Inr o inj
```

Note that these type class instances overlap, going beyond Haskell 98; nevertheless, as Swierstra explains, provided that sums are associated to the right this should not cause any problems.

Now we can define smart constructors that inject in this ‘obvious’ way:

```
identity11 :: (Identity11 <: f) => Width -> Fix f
identity11 w = In (inj (Identity11 w))

fan11 :: (Fan11 <: f) => Width -> Fix f
fan11 w = In (inj (Fan11 w))

above11 :: (Above11 <: f) => Fix f -> Fix f -> Fix f
above11 x y = In (inj (Above11 x y))

beside11 :: (Beside11 <: f) => Fix f -> Fix f -> Fix f
beside11 x y = In (inj (Beside11 x y))

stretch11 :: (Stretch11 <: f) => [Width] -> Fix f -> Fix f
stretch11 ws x = In (inj (Stretch11 ws x))
```

and the mess of injections can be inferred instead:

```
stretchfan :: (Fan11 <: f, Stretch11 <: f) => Fix f
stretchfan = stretch11 [3,2,3] (fan11 3)
```

Crucially, this technique also leaves the precise choice of grammar open; all that is required is for the grammar to provide fan and stretch constructors, and we can capture that dependence in the flexible declared type for *stretchfan*.

Interpretations can be similarly modularized. Of course, we expect them to be folds:

```
fold :: Functor f => (f a -> a) -> Fix f -> a
fold h (In x) = h (fmap (fold h) x)
```

In order to accommodate open datatypes, we define interpretations in pieces. We declare a type class of those constructors supporting a given interpretation:

```
class Functor f => WidthAlg f where
  widthAlg11 :: f Width -> Width
```

Interpretations lift through sums in the obvious way:

```
instance (WidthAlg f, WidthAlg g) => WidthAlg (f :+: g) where
  widthAlg11 (Inl x) = widthAlg11 x
  widthAlg11 (Inr y) = widthAlg11 y
```

Then we provide instances for each of the relevant constructors. For example, if we only ever wanted to compute the width of circuits expressed in terms of the fan and stretch constructors, we need only define those two instances:

```
instance WidthAlg Fan11 where
  widthAlg11 (Fan11 w) = w

instance WidthAlg Stretch11 where
  widthAlg11 (Stretch11 ws x) = sum ws
```

For example, this width function works for the flexibly typed circuit *stretchfan* above:

```
width11 :: WidthAlg f => Fix f -> Width
width11 = fold widthAlg11
```

—although the circuit does need to be given a specific type first:

```
> width11 (stretchfan :: Circuit11)
8
```

These algebra fragments together constitute the essence of an implicitly parametrized shallow embedding.

But the main benefit of the *à la carte* approach is that it is easy to add new constructors. We just need to add the datatype constructor as a functor, and provide a smart constructor:

```
data RStretch11 c = RStretch11 [Size] c deriving Functor
rstretch11 :: (RStretch11 <: f) => [Width] -> Fix f -> Fix f
rstretch11 ws x = In (inj (RStretch11 ws x))
```

Now the circuit in Figure 4 can be expressed using right stretch instead of left stretch:

```
rstretchfan :: (Identity11 <: f, Fan11 <: f, Beside11 <: f,
  RStretch11 <: f) => Fix f
rstretchfan = beside11 (identity11 2)
  (rstretch11 [2,3,1] (fan11 3))
```

When adding new constructors such as *RStretch₁₁*, it is tempting to provide an instance for each of the interpretations of interest, such as *WidthAlg*. However, this is an unnecessary duplication of effort when *rstretch₁₁* can itself be simulated out of existing components. We might instead write a function that *handles* the *RStretch₁₁* constructor:

```
handle :: (Stretch11 <: f, Beside11 <: f, Identity11 <: f) =>
  Fix (RStretch11 <: f) -> Fix f
handle (In (Inl (RStretch11 ws c))) =
  stretch11 (1 : ws') (handle c) 'identity11 w
  where (ws', w) = (init ws, last ws - 1)
handle (In (Inr other)) = In (fmap handle other)
```

Here, we recursively translate all instances of *RStretch₁₁* into other constructors. This technique is at the heart of the effects and handlers approach [21], although the setting there uses the free monad rather than *Fix*. With this in place, we can first handle all of the *RStretch₁₁* constructors before passing the result on to an interpretation function such as *width₁₁* that need not deal with *RStretch₁₁*s. This method of interpreting only a core fragment of syntax might not be optimally efficient, but of course we still leave open the possibility of providing a specialized instance if that is an issue.

5. Discussion

The essential observation made here—that *shallow embeddings correspond to the algebras of folds over the abstract syntax captured by a deep embedding*—is surely not new. For example, it was probably known to Reynolds [25], who contrasted deep embeddings (‘user defined types’) and shallow (‘procedural data structures’), and observed that the former were free algebras; but he didn’t explicitly discuss anything corresponding to folds.

It is also implicit in the *finally tagless* approach [8], which uses a shallow embedding and observes that ‘this representation makes it trivial to implement a primitive recursive function over

object terms’, providing an interface that such functions should implement; but this comment is made rather in passing, and their focus is mainly on staging and partial evaluation. The observation is more explicit in Kiselyov’s lecture notes on the finally tagless approach [20], which go into more detail on compositionality; he makes the connection to “denotational semantics, which is required to be compositional”, and observes that “making context explicit turns seemingly non-compositional operations compositional”. The finally tagless approach also covers DSLs with binding constructs, which we have ignored here.

Neither is it a new observation that algebraic datatypes (such as *Circuit*₄) and their Böhm–Berarducci encodings (such as *Circuit*₇) are equivalent. And of course, none of this is specific in any way to the *Circuit* DSL; a datatype-generic version of the story can be told, by abstracting away from the shape functor *CircuitF*—the reader may enjoy working out the details.

Nevertheless, the observation that shallow embeddings correspond to the algebras of folds over deep embeddings seems not to be widely appreciated; at least, we have been unable to find an explicit statement to this effect, either in the DSL literature or elsewhere. And it makes a nice application of folds: many results about folds evidently have interesting statements about shallow embeddings as corollaries. The three generalizations of folds (banana split, mutumorphisms, and accumulating parameters) exploited in Section 4 are all special cases of *adjoint fold* [15, 16]; perhaps other adjoint folds yield other interesting insights about shallow embeddings?

Acknowledgements

This paper arose from ideas discussed at the Summer School on Domain Specific Languages in Cluj-Napoca in July 2013 [11]; JG thanks the organizers for the invitation to lecture there. José Pedro Magalhães, Ralf Hinze, Jacques Carette, James McKinnin, and the anonymous reviewers all made helpful comments, and Oleg Kiselyov gave many constructive criticisms and much inspiration, for which we are very grateful. This work has been funded by EPSRC grant number EP/J010995/1, on Unifying Theories of Generic Programming.

References

- [1] Robert Atkey, Sam Lindley, and Jeremy Yallop. Unembedding domain-specific languages. In *Haskell Symposium*, pages 37–48. ACM, 2009.
- [2] Lennart Augustsson. Making EDSLs fly. In *TechMesh*, London, December 2012. Video at <http://vimeo.com/73223479>.
- [3] Richard S. Bird. The promotion and accumulation strategies in transformational programming. *ACM Transactions on Programming Languages and Systems*, 6(4):487–504, October 1984. Addendum in *TOPLAS* 7(3):490–492, July 1985.
- [4] Corrado Böhm and Alessandro Berarducci. Automatic synthesis of typed λ -programs on term algebras. *Theoretical Computer Science*, 39:135–154, 1985.
- [5] Richard Boulton. Personal communication, 10th February 2014.
- [6] Richard Boulton, Andrew Gordon, Mike Gordon, John Harrison, John Herbert, and John Van Tassel. Experience with embedding hardware description languages in HOL. In Victoria Stavridou, Thomas F. Melham, and Raymond T. Boute, editors, *IFIP TC10/WG 10.2 International Conference on Theorem Provers in Circuit Design: Theory, Practice and Experience*, pages 129–156. North-Holland/Elsevier, 1992.
- [7] Richard P. Brent and Hsiang-Tsung Kung. The chip complexity of binary arithmetic. In *Symposium on Theory of Computing*, pages 190–200. ACM, 1980.
- [8] Jacques Carette, Oleg Kiselyov, and Chung-chieh Shan. Finally tagless, partially evaluated: Tagless staged interpreters for simpler typed languages. *Journal of Functional Programming*, 19(5):509–543, 2009.
- [9] Maarten M. Fokkinga. Tupling and mutumorphisms. *The Squiggolist*, 1(4):81–82, June 1990.
- [10] Martin Fowler. *Domain-Specific Languages*. Addison-Wesley, 2011.
- [11] Jeremy Gibbons. Functional programming for domain-specific languages. In Viktória Zsóka, editor, *Central European Functional Programming Summer School*, volume 8606 of *Lecture Notes in Computer Science*, pages 1–27. Springer, 2014. To appear.
- [12] Jeremy Gibbons and Geraint Jones. The under-appreciated unfold. In *International Conference on Functional Programming*, pages 273–279, Baltimore, Maryland, September 1998.
- [13] Ralf Hinze. An algebra of scans. In *Mathematics of Program Construction*, volume 3125 of *Lecture Notes in Computer Science*, pages 186–210. Springer, 2004.
- [14] Ralf Hinze. Church numerals, twice! *Journal of Functional Programming*, 15(1), 2005.
- [15] Ralf Hinze. Adjoint folds and unfolds: An extended study. *Science of Computer Programming*, 78(11):2108–2159, 2013.
- [16] Ralf Hinze, Nicolas Wu, and Jeremy Gibbons. Unifying structured recursion schemes. In *International Conference on Functional Programming*, pages 209–220, Boston, Massachusetts, September 2013.
- [17] Zhenjiang Hu, Hideya Iwasaki, and Masato Takeichi. Formal derivation of efficient parallel programs by construction of list homomorphisms. *ACM Transactions on Programming Languages and Systems*, 19(3):444–461, 1997.
- [18] Paul Hudak. Building domain-specific embedded languages. *ACM Computing Surveys*, 28(4), 1996.
- [19] Anne Kaldewaij. *Programming: The Derivation of Algorithms*. Prentice Hall, 1990.
- [20] Oleg Kiselyov. Typed tagless final interpreters. In Jeremy Gibbons, editor, *Generic and Indexed Programming*, volume 7470 of *Lecture Notes in Computer Science*, pages 130–174. Springer, 2012.
- [21] Oleg Kiselyov, Amr Sabry, and Cameron Swords. Extensible effects: An alternative to monad transformers. In *Haskell Symposium*, pages 59–70. ACM, 2013.
- [22] Shriram Krishnamurthi, Matthias Felleisen, and Daniel P. Friedman. Synthesizing object-oriented and functional design to promote re-use. In *European Conference on Object-Oriented Programming*, volume 1445 of *Lecture Notes in Computer Science*, pages 91–113. Springer, 1998.
- [23] Ehud Lamm. CUFP write-up. Blog post, <http://lambda-the-ultimate.org/node/2572>, December 2007.
- [24] Ralf Lämmel, Joost Visser, and Jan Kort. Dealing with large bananas. In Johan Jeuring, editor, *Workshop on Generic Programming*, volume Technical Report UU-CS-2000-19. Universiteit Utrecht, 2000.
- [25] John Reynolds. User-defined types and procedural data structures as complementary approaches to data abstraction. In Stephen A. Schuman, editor, *New Directions in Algorithmic Languages*, pages 157–168, 1975.
- [26] Josef Svenningsson and Emil Axelsson. Combining deep and shallow embedding for embedded domain-specific languages. In *Trends in Functional Programming 2012*, volume 7829 of *Lecture Notes in Computer Science*, pages 21–36, 2013.
- [27] Wouter Swierstra. Datatypes à la carte. *Journal of Functional Programming*, 18(4):423–436, 2008.
- [28] Philip Wadler. Theorems for free! In *Functional Programming Languages and Computer Architecture*, pages 347–359. ACM, 1989.
- [29] Philip Wadler. The expression problem. Java Genericity Mailing list, November 1998. <http://homepages.inf.ed.ac.uk/wadler/papers/expression/expression.txt>.
- [30] Philip Wadler and Stephen Blott. How to make ad-hoc polymorphism less ad hoc. In *Principles of Programming Languages*, pages 60–76. ACM, 1989.